

Does simultaneous solution matter for stellar evolution codes?

Richard J. Stancliffe*

Institute of Astronomy, The Observatories, Madingley Road, Cambridge CB3 0HA, U.K.

Accepted 0000 December 00. Received 0000 December 00; in original form 0000 October 00

ABSTRACT

A version of the STARS stellar evolution code has been developed that uses a non-simultaneous solution of the equations of stellar structure and evolution. In all other respects it is identical to the normal, fully simultaneous version. It is therefore possible to test the dependence of the solution on how the equations are solved. Two cases are investigated: a $5 M_{\odot}$ and a $3 M_{\odot}$ star, both of metallicity $Z=0.02$. Prior to the asymptotic giant branch, the models are almost identical. However once thermal pulses start, the two methods of solution yield diverging results with the non-simultaneous technique predicting longer interpulse periods. This is traced to difficulties associated with hydrogen burning caused by the use of a moving mesh. It is shown that, with careful control of the temporal resolution, the results of the simultaneous technique can be recovered.

Key words: methods: numerical, stars: evolution, stars: AGB and post-AGB

1 INTRODUCTION

Stellar evolution is based upon the solution of the equations of stellar structure, together with equations governing the mixing and burning of certain important isotopes. For a spherically symmetric star that is non-rotating, these equations are (see e.g. Kippenhahn & Weigert 1990):

- the equation of hydrostatic equilibrium,

$$\frac{dP}{dm} = -\frac{Gm}{4\pi r^4}, \quad (1)$$

where P, r and m are pressure, radius and the mass contained within a spherical shell of radius r respectively.

- The equation of mass conservation,

$$\frac{dr}{dm} = \frac{1}{4\pi r^2 \rho}, \quad (2)$$

where ρ is density.

- The equation of energy conservation,

$$\frac{dL}{dm} = \epsilon, \quad (3)$$

where L is luminosity and ϵ is the energy generation rate including nuclear energy generation, energy from gravitational sources and energy losses from neutrino emission.

- The equation of energy transport, which may be expressed as

$$\frac{d \ln T}{dm} = -\nabla \frac{d \ln P}{dm} \quad (4)$$

where the form of ∇ depends on whether the region of the star is radiative or convective.

Assuming mixing is treated as a diffusive process¹, the equation governing mixing and burning for each isotope X_i is given by:

$$\frac{d}{dm} \left(\sigma \frac{dX_i}{dm} \right) = \frac{dX_i}{dt} + R_i - S_i, \quad (5)$$

where σ is the diffusion coefficient, R_i is the rate at which the species i is being burnt by nuclear reactions and S_i is the rate at which it is being produced by nuclear reactions (Eggleton 1972).

While these equations may be common to all evolution codes, the method employed to solve them varies. Three possible ways in which current codes can solve these equations may be identified. Before describing these methods it is necessary to define some terminology. A *timestep* is defined as the act of moving from a model at time t to one at time $t + \Delta t$. To produce a model at a new timestep using the relaxation method it is necessary to make *iterations* on the solution. When the changes to the current solution are sufficiently small the model is said to have *converged*. Bearing in mind these definitions, three approaches to solving the equations may be defined: non-simultaneous, partially simultaneous and fully simultaneous.

The non-simultaneous approach involves converging a solution for the structure for a given timestep. This newly-

* E-mail: rs@ast.cam.ac.uk

¹ Not all codes treat mixing in this way.

computed structure is then used to calculate the mixing and burning for that timestep. Iterations are made separately on the structure and the chemistry (mixing and burning). Examples of codes that employ this method include those of Straniero et al. (2000) and Herwig (2000). The partially simultaneous approach involves solving for the structure equations for an iteration, then performing an iteration on the mixing/burning. One continues alternating between structure and chemistry iterations until the next timestep is converged. Such an approach is used in the Mount Stromlo Stellar Structure Program (MSSSP) employed by Karakas et al. (2002), for example. Finally, the fully simultaneous approach involves solving all the equations together at each individual iteration of each timestep. This is the method employed by Eggleton (1971) in the original development of the STARS code and continues to be used in its current incarnations.

Current evolution codes predict different results for computations of the thermally pulsing asymptotic giant branch (Lugaro et al. 2003; Stancliffe et al. 2004). The reason for these discrepancies is uncertain. This paper addresses the effect the method of solution has on the results obtained. While it is thought that simultaneous solution will not influence the results of evolution calculations (Herwig 2005), the effects have not been tested until now.

2 THE STARS CODE

The stellar evolution code STARS, originally developed by Eggleton (1971) and most recently updated by Pols et al. (1995) was used in this study. Derivatives of the original Eggleton code are unique in that they are the only codes to solve the equation of stellar structure and evolution in a fully simultaneous manner.

One would expect that if any difference exists between the methods of solution, it is likely to be greatest between the non-simultaneous and the fully simultaneous approaches. Therefore a version of the STARS code that employs a non-simultaneous method of solution has been developed. To do this, use has been made of a pair of subroutines designed to follow the evolution of minor isotopes (see Stancliffe et al. 2005; Stancliffe 2005, for further details). These routines take the structure computed by the main routines and use it to compute the burning and mixing of material which is energetically unimportant.

The version of the code employing a non-simultaneous method of solution works as follows. Equations 1-4 are solved for using the normal evolution routines, with the mesh points being moved to their most appropriate locations by the mesh spacing function. It is therefore necessary to solve for the change in composition due to advection. Once a solution has been converged for the structure and the composition change due to mesh movement, equation 5 is then solved for each of the five important isotopes (^1H , ^4He , ^{12}C , ^{14}N , ^{16}O) using the minor element subroutines, using the structure that has just been converged. The mesh spacing remains the same as for the evolution step and the diffusive mixing coefficient is taken from the converged structure. Thus it is possible to convert the method of solution employed by the STARS code into a non-simultaneous one by removing the equations of burning and mixing from the main code and placing them in these subroutines.

By changing the method of solution used by a code, we have two codes that are identical in all respects except for the one we are interested in. It is therefore possible to compare the effect that non-simultaneous solution *alone* has on the results obtained. This sort of comparison would be extremely difficult to do across different codes as different codes employ slightly different reaction rates, equations of state, mixing prescriptions, etc. and these would introduce uncertainties in to the source of any difference obtained. In addition there are different ways to write the difference equations used, different ways of averaging certain quantities (see e.g. Pols & Tout 2001) and different choices for defining variables at the centre or edges of shells. All these could affect the results obtained. By creating a version of the STARS code that uses a non-simultaneous solution these problems are completely bypassed. In addition, the timestep control can be kept the same for both codes.

The STARS code attempts to choose the most appropriate timestep size based on the changes to the variables of the previous model required to produce the current model. A sum of the absolute values of these changes (excluding those made to the luminosity) is made over all variables and over all meshpoints producing a single numerical value

$$d = \sum_i \sum_k |\Delta x_{i,k}|,$$

where $\Delta x_{i,k}$ are the values of the changes made to each variable i at a given meshpoint k . This is then compared to a preset optimum value d_{opt} . If d_{opt}/d is greater than one the timestep is increased by this fraction or 1.2, whichever is smaller. If d_{opt}/d is less than one the timestep is reduced by this fraction or 0.8, whichever is larger. In this way the most appropriate timestep is chosen for the next model. The value of d is dominated by the temperature and the degeneracy; the remaining variables make only minor contributions to the timestep control.

3 RESULTS

To compare the differences between the two codes a $5 M_{\odot}$ star of initial metallicity $Z=0.02$ was evolved from the pre-main sequence to the asymptotic giant branch without mass loss or convective overshooting using both simultaneous and non-simultaneous approaches. The initial model had 499 mesh points. Once second dredge-up was over, the models were remeshed with 1999 mesh points² and the AGB specific mesh spacing function of Stancliffe, Tout & Pols (2004) was employed in order to ensure proper spatial resolution of the important features of AGB stars.

Throughout the pre-main sequence, main sequence and red giant branch the evolution computed by the two methods is indistinguishable. At core helium burning the non-simultaneous solution displays a blue loop that extends toward the blue more than the simultaneous solution but this deviation is less than a tenth of a percent and should not be considered significant.

As the asymptotic giant branch is ascended the simultaneous solution grows a larger core than the non-simultaneous

² This is twice the number normally employed and was used in case making the code non-simultaneous required additional spatial resolution.

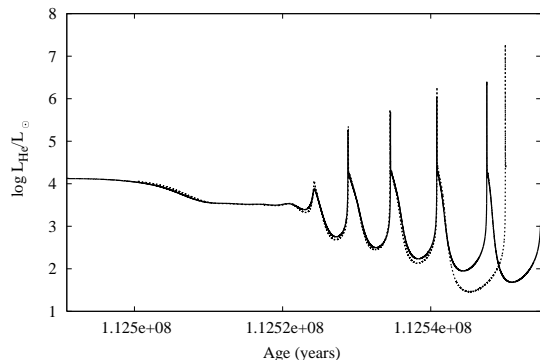


Figure 1. Evolution of the helium luminosity as a function of age. The simultaneous solution is shown as a solid line; the non-simultaneous solution is the dashed line. The non-simultaneous solution has been shifted forward by 4.44×10^4 yr to make the first thermal pulses coincident.

one. Prior to second dredge-up the simultaneous solution has a hydrogen-exhausted core that is about 0.3% larger than the core of the non-simultaneous solution. The simultaneous solution then undergoes deeper second dredge-up than the non-simultaneous solution, resulting in it having a core that is 0.6% less massive. This difference is about an order of magnitude smaller than the difference between the simultaneous solution and the comparable model of Karakas et al. (2002), computed with a code employing a partially simultaneous method of solution. It therefore seems unlikely that the differences in details such as the core mass at first thermal pulse are due to the method of solution.

At the beginning of the thermally pulsing asymptotic giant branch (TP-AGB) the two models retain their similarities for a few thermal pulses. A rapid divergence in behaviour then sets in. This is demonstrated in the evolution of the helium luminosities for both models as shown in Figure 1. The first three thermal pulses all have approximately similar peak helium luminosities. However, the interpulse period following the third thermal pulse is much longer for the non-simultaneous solution. This then results in the fourth thermal pulse of that model being much stronger.

How much of the deviation is due to the fact that the two models have slightly different core mass at the beginning of the TP-AGB? While the initial difference is minimal it is not inconceivable that a small difference in initial conditions can lead to a larger difference later in the evolution. To assess the impact of just the difference in the method of solution, a starting model for a run using non-simultaneous solution was taken from the simultaneous solution run, just after second dredge-up. The evolution of the helium luminosity compared to the evolution from the simultaneous solution run is shown in Figure 2. Again, the initial pulses are similar but are seen to diverge after a couple of pulses. It would be desirable to continue the sequence to compare the behaviour over a longer run of thermal pulses but a numerical instability that could not be overcome in the non-simultaneous solution sequence meant this could not be done.

It therefore seems that simultaneous and non-simultaneous methods of solution do not yield the same re-

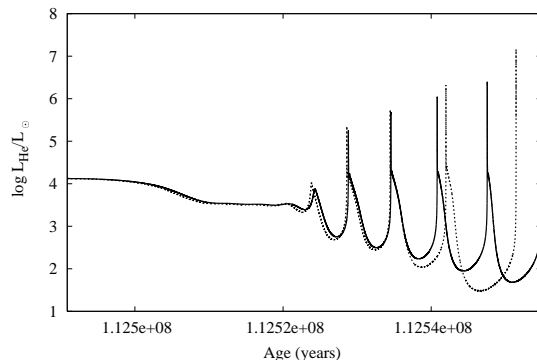


Figure 2. Evolution of the helium luminosity as a function of age. The simultaneous solution is shown as a solid line; the non-simultaneous solution is the dashed line. Both sequences have been given the same initial model which was taken from just after second dredge-up.

sults, despite the same timestep control routines being used in both cases.

3.1 $3 M_{\odot}$ models

Having noted that the two solution methods give different results, it is important to test whether these results extended to lower masses. Therefore $3 M_{\odot}$, $Z=0.02$ models were also produced using both codes. The details of the thermal pulse runs can be found in Tables 1 and 2.

As with the $5 M_{\odot}$ models, the $3 M_{\odot}$ also display different behaviour, despite entering the TP-AGB with core masses within less than a percent of one another. The main difference again seems to be in the interpulse periods which are again larger in the case of the non-simultaneous solution. The longer interpulse period also means that there is greater core growth in the non-simultaneous solution. The non-simultaneous solution also gives more violent thermal pulses. This is probably related to the longer interpulse period (see the discussion in Stancliffe et al. 2004).

In terms of the onset of third dredge-up and its depth, both models give similar results. In both cases TDUP sets in when the H-exhausted core mass exceeds $0.585 M_{\odot}$ and this happens on the sixth pulse in both models. Over the few thermal pulses with TDUP both models give comparable dredge-up efficiencies, with the simultaneous solution giving slightly lower efficiencies (which may be due to the slightly lower pulse strength). Curiously, while the simultaneous solution gives slightly less TDUP its C/O ratio becomes higher than in the case of the non-simultaneous solution. This occurs because the envelope in the case of the non-simultaneous solution goes deeper into the star during TDUP. This results in the carbon being mixed in with a greater amount of envelope material and hence becoming more diluted.

3.2 Testing TDUP

Does the use of simultaneous solution affect the amount of TDUP predicted? It is difficult to assess the difference

TP	M_{H} (M_{\odot})	τ_{p} (10^4 yr)	$\log (L_{\text{He}}^{\text{max}} / L_{\odot})$	ΔM_{H} (M_{\odot})	ΔM_{DUP} (M_{\odot})	λ	C/O
1	0.55267	...	4.36792	0.01628	0.315
2	0.56894	7.60	6.27639	0.00440	0.315
3	0.57333	7.78	5.56767	0.00297	0.315
4	0.57629	8.89	6.65021	0.00580	0.315
5	0.58208	9.33	6.62750	0.00574	0.315
6	0.58781	8.81	6.80665	0.00621	0.00047	0.076	0.317
7	0.59355	8.22	6.90432	0.00633	0.00113	0.179	0.330
8	0.59875	7.76	7.03324	0.00660	0.00190	0.288	0.356
9	0.60345	7.50	7.16711	0.00701	0.00266	0.379	0.393

Table 1. Details of the $3 M_{\odot}$ model computed using simultaneous solution. The data are TP – the thermal pulse number, M_{H} – the hydrogen free core mass, τ_{p} – the interpulse period, $L_{\text{He}}^{\text{max}}$ – the peak luminosity from helium burning, ΔM_{H} – the hydrogen free core mass growth during the interpulse, ΔM_{DUP} – the mass of material dredged up, $\lambda = \Delta M_{\text{DUP}} / \Delta M_{\text{H}}$ – the dredge-up efficiency and C/O – the surface carbon-to-oxygen ratio by number.

TP	M_{H} (M_{\odot})	τ_{p} (10^4 yr)	$\log (L_{\text{He}}^{\text{max}} / L_{\odot})$	ΔM_{H} (M_{\odot})	ΔM_{DUP} (M_{\odot})	λ	C/O
1	0.54981	...	5.64095	0.01704	0.315
2	0.56684	6.87	5.33562	0.00241	0.315
3	0.56925	8.99	6.45104	0.00517	0.315
4	0.57441	10.99	6.65133	0.00591	0.315
5	0.58032	10.59	6.87245	0.00642	0.315
6	0.58666	10.01	6.96026	0.00663	0.00064	0.097	0.318
7	0.59265	9.52	7.12672	0.00707	0.00142	0.201	0.328
8	0.59830	9.17	7.28263	0.00755	0.00222	0.294	0.347

Table 2. Details of the $3 M_{\odot}$ model computed using non-simultaneous solution. The headings are the same as in Table 1.

caused by using a simultaneous solution from the preceding runs because of the relationship between longer interpulse periods, more violent thermal pulses and deeper third dredge-up.

A sequence of models is therefore generated using both the simultaneous and non-simultaneous methods of solution, taking a starting model from just prior to the onset of TDUP after the 10th thermal pulse of the simultaneous-solution run for the $3 M_{\odot}$ star described above. The results are shown in Figure 3. The effect is quite dramatic: the non-simultaneous solution gives much less TDUP. For the simultaneous solution 0.00366 M_{\odot} of intershell material is dredged-up while for the non-simultaneous solution only 0.00283 M_{\odot} is dredged-up. This corresponds to a dredge-up efficiency, λ , of 0.493 for the simultaneous solution versus 0.381 for the non-simultaneous solution.

4 DISCUSSION

From the sample runs provided above it appears that there does exist a difference between the simultaneous and non-simultaneous solution of the equations of stellar structure and evolution. However, it appears that the important quantities of the core mass at first thermal pulse and the core mass at which third dredge-up occurs are not affected substantially by the use of a simultaneous or non-simultaneous solution as they differ by less than a percent in both these runs. By contrast, the difference between the core mass at first thermal pulse in the run using simultaneous solution is around 7% smaller than the corresponding model of the

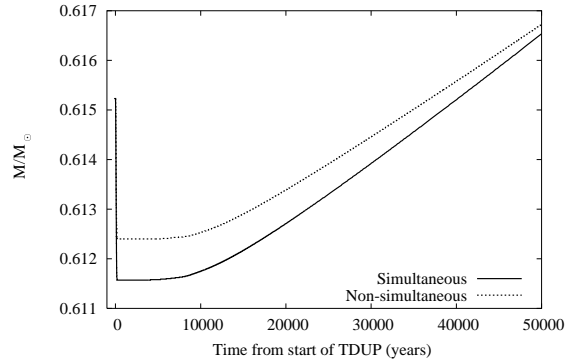


Figure 3. Plot of the depth of TDUP following the 10th thermal pulse of the $3 M_{\odot}$ sequence, using a starting model taken from the simultaneous-solution sequence. The run using simultaneous solution is the solid line; the dashed line is the non-simultaneous solution.

code of Karakas et al. (2002), which uses a partially simultaneous method of solution. This would suggest that the differences that occur between the predictions are not due to the method of solution used.

What is the origin of the divergence in the model predictions presented above? The difference is first noticed in the interpulse period, which is when the H-burning shell is active. This is a logical place to start looking for a reason for the divergence. Figure 4 shows the evolution of the H-burning luminosity for a $1.5 M_{\odot}$ $Z=0.02$ model, evolved

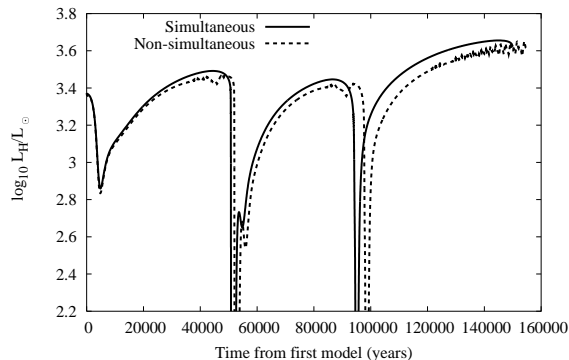


Figure 4. Plot of the evolution of the hydrogen-burning luminosity as a function of time for a $1.5 M_{\odot}$ model. The simultaneous solution is the solid line; the dashed line is the non-simultaneous solution. Note the erratic behaviour of the non-simultaneous solution prior to the drop in the H-luminosity at each thermal pulse.

from the same starting model with both codes. The non-simultaneous solution shows erratic behaviour in its hydrogen luminosity while the simultaneous solution is smooth as would be expected.

The erratic behaviour of the H-burning luminosity in the non-simultaneous solution points to a resolution problem. By forcing the code employing the non-simultaneous solution to use smaller timesteps the simultaneous solution can be approached (see Figure 5). To obtain the high temporal resolution run the timestep has been reduced by at least an order of magnitude relative to the simultaneous solution run. The fluctuations in the timestep³ around just prior to the thermal pulse suggest that there are still problems with the model even at these lower timesteps. A high spatial resolution model consisting of 4999 mesh points and using the same timestep control as the original non-simultaneous solution run gives the same erratic behaviour in the H-luminosity as the original run. This shows that the resolution problem is a temporal issue rather than a spatial one.

It therefore appears that a code employing a non-simultaneous method of solution can reproduce the results of one using a simultaneous solution, providing that more care is taken over the size of the timestep used. The bottom panel of Figure 5 also shows that even at low timesteps the non-simultaneous solution still has difficulty with hydrogen burning as shown by the variation in timestep size around 40,000 years from the initial model. Similar behaviour is found when the issue of TDUP is re-examined in this light. The non-simultaneous solution is seen to give comparable dredge-up to the non-simultaneous solution when the timestep size is reduced by over an order of magnitude.

4.1 Lagrangian versus non-Lagrangian codes

One caveat remains: the method used above to produce a code using a non-simultaneous method of solution does not produce one that is identical to the way most stellar evolution codes operate. Most codes in use are Lagrangian, i.e.

³ They are caused by a timestep failing to converge. In such a case, the code tries again using a smaller timestep size.

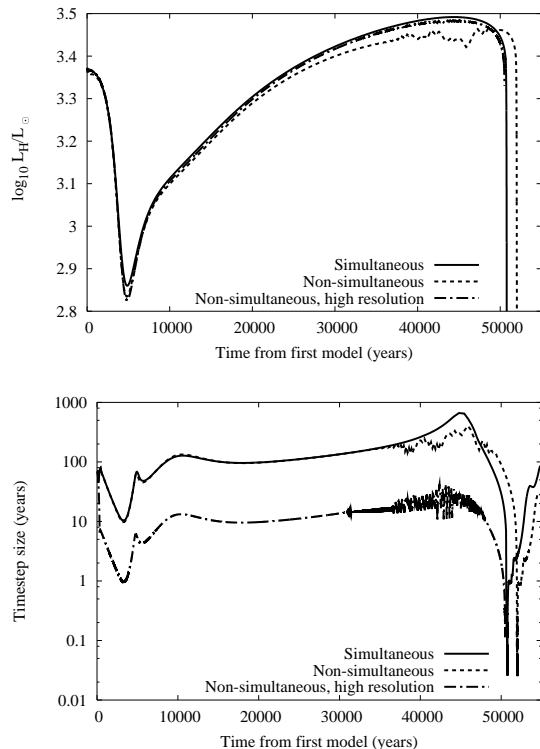


Figure 5. Top panel: Plot of the evolution of the H-burning luminosity as a function of time for the $1.5 M_{\odot}$ model. The solid line is the simultaneous solution, the dotted line is the non-simultaneous solution and the dot-dashed line is the non-simultaneous solution with higher temporal resolution. Bottom panel: Variation in timestep size as a function of time for the models in the top panel.

the mesh points are placed at locations with fixed mass coordinates. This is not true of the STARS code. It was designed to use a non-Lagrangian, non-Eulerian mesh with a fixed number of mesh points; the mesh points are allowed to move in order to follow physically important features (Eggleton 1971). Eggleton's adaptive mesh (hereafter referred to as an Eggletonian mesh) has been retained in the version of the code employing a non-simultaneous method of solution.

It is possible that the combination of a non-simultaneous solution with an Eggletonian mesh is responsible for the instabilities outlined above. When the mesh moves during a structural step, the compositions are also changed to account for the movement. Where there are sharp gradients in composition this can cause problems. Suppose a profile with a sharp discontinuity is set up over a series of mesh points as depicted by the solid line in Figure 6, with the black circles representing the location of the mesh points. If the mesh points move inward in mass, as indicated, to the locations marked by the grey circles, the profile is modified to that indicated by the dashed line. This is because when a mesh point moves the code works out the value of the abundance at the new location from the values from the old mesh point locations surrounding the new one. This inevitably leads to the sharp profile being smeared out. It is an unfortunate and very much undesirable consequence of using an adaptive mesh.

With this numerical diffusion, extra fuel can end up

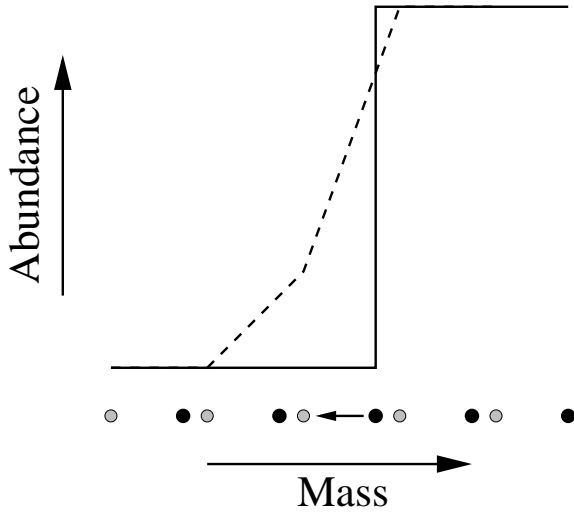


Figure 6. A schematic depiction of why numerical diffusion occurs. A sharp abundance profile (solid line) is set up over a few mesh points (black dots). If the mesh points then move in mass to the positions indicated by the grey circles, the abundance profile is modified to that indicated by the dashed line.

where it physically should not be. This results in the burning rate being increased. When the equations are solved, feedback to the structure can result and if it is strong enough it may create an instability. In a code employing simultaneous solution, this feedback is dealt with during the iterations required to converge a model so the problem does not arise. In a code employing a non-simultaneous solution, if the mesh moves slowly, or the timestep is low enough then the feedback between the structure and composition can be dealt with. The question is: how low must the timestep be for this to happen? As a naïve expectation, the timescale should be comparable to the time it would take for a given mesh point to traverse the mass difference between adjacent mesh points, i.e.:

$$\tau_{\text{drift}} = \left| \frac{m_k - m_{k+1}}{\dot{m}_k} \right| \quad (6)$$

where m is the mass of the k th mesh point and \dot{m}_k is the rate of change of mass of the k th mesh point due to the movement of the mesh (any given mesh point may move inward or outward in mass, depending on where the mesh spacing function thinks it should be). If the timestep of the code is smaller than this then changes should be properly resolved. In the above $1.5 M_{\odot}$ model, the drift timescale τ_{drift} in the hydrogen burning shell is a little under 100 years and the timestep size is of this order. An instability is seen in the hydrogen burning luminosity. When smaller timesteps are used⁴, the timestep size is much less than the drift timescale and no instability is detected.

This problem would be avoided by using a Lagrangian mesh. From Equation 6 it can be seen that as \dot{m} tends to zero (i.e. the mesh tends towards being Lagrangian), τ_{drift} tends to infinity and the code will always have a timestep much smaller than the drift time. This could be tested by creating a Lagrangian code from the non-simultaneous code

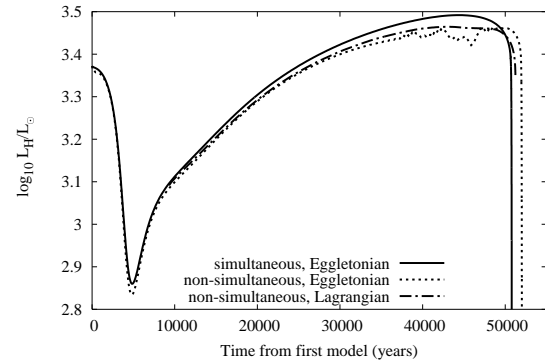


Figure 7. Plot of the evolution of the hydrogen-burning luminosity as a function of time for the $1.5 M_{\odot}$ model. The solid line represents the simultaneous solution which has an Eggletonian mesh, the dotted line represents the non-simultaneous code with the Eggletonian mesh and the dot-dashed line is the non-simultaneous code with a Lagrangian mesh near its surface and an Eggletonian mesh elsewhere. Note the absence of instabilities in the latter track.

used here but this would be a major undertaking and is beyond the scope of this work.

However, it is possible to make a limited test of the situation. Using the viscous mesh technique of Stancliffe et al. (2004) it is possible to fix the mass locations of some of the mesh points. In this way, the regions of the mesh that are believed to be the problem (i.e. those mesh points from the surface to below the hydrogen burning shell) can be made to behave as Lagrangian mesh points. Unfortunately, when the mesh is locked it affects the timestep control mechanism typically producing lower timesteps than the unlocked mesh. Care must be taken to ensure that the effects of using a low timestep and the effects of using a Lagrangian mesh are not confused.

The $1.5 M_{\odot}$ model was re-run with the 700 mesh points closest to the surface being fixed in their mass co-ordinates. The base of the H-burning shell initially lies at around 500 mesh points from the surface so the problematic region is within the Lagrangian part of the mesh. Figure 7 shows the result. The H-luminosity evolves smoothly, though it does not replicate the behaviour of the simultaneous solution using the Eggletonian mesh. This is believed to be due to a lack of a resolution around the shell as a similar result is obtained if the upper 700 mesh points of the code employing the simultaneous solution are fixed.

It seems that if a non-simultaneous method of solution is to be used then a Lagrangian mesh is an advantage, as numerical diffusion in an Eggletonian mesh can lead to instability. However, the use of a fully simultaneous method of solution seems to avoid these problems. Further work on the effects of numerical diffusion within the framework of the STARS code needs to be done.

5 CONCLUSIONS

The dependence of the solution obtained on the method by which the equations of stellar structure and evolution are solved has been investigated for low and intermediate-mass

⁴ This will reduce the magnitude of \dot{m} and hence increase τ_{drift} .

stars. No difference was found prior to the AGB and the method of solution does not account for the differences seen in, for example, the core mass at first thermal pulse as computed by different evolution codes.

Along the TP-AGB the two methods give diverging results, with the non-simultaneous method giving longer interpulse periods. This leads to more violent thermal pulses. The difference in interpulse period was tracked to an instability in the hydrogen burning shell of the non-simultaneous solution. This is due to numerical diffusion caused by the use of an Eggletonian mesh. By using smaller timesteps, or by using a Lagrangian mesh, the instability could be overcome.

The problems arising from the use of an Eggletonian mesh cloud the issue of a direct comparison between the non-simultaneous and simultaneous methods of solution. The discrepancies observed above are shown to be related to the use of an Eggletonian mesh together with a non-simultaneous method of solution. There is no evidence for an inherent problem with using a non-simultaneous method of solution. Further work needs to be done to assess the effect of using an Eggletonian mesh.

6 ACKNOWLEDGEMENTS

The author is grateful to the referee, John Lattanzio, for his careful reading of the manuscript and his many helpful suggestions that have added to the clarity of this work. RJS thanks Pierre Lesaffre for many useful discussions greatly improving mental, if not numerical, stability. He also thanks Churchill College for his fellowship.

REFERENCES

- Eggleton P. P., 1971, MNRAS, 151, 351
- Eggleton P. P., 1972, MNRAS, 156, 361
- Herwig F., 2000, A&A, 360, 952
- Herwig F., 2005, ARA&A, 43, 435
- Karakas A. I., Lattanzio J. C., Pols O. R., 2002, Pub. Astron. Soc. Aust., 19, 515
- Kippenhahn R., Weigert A., 1990, Stellar Structure and Evolution. Stellar Structure and Evolution, XVI, 468 pp. 192 figs.. Springer-Verlag Berlin Heidelberg New York. Also Astronomy and Astrophysics Library
- Lugaro M., Herwig F., Lattanzio J. C., Gallino R., Straniero O., 2003, ApJ, 586, 1305
- Pols O. R., Tout C. A., 2001, Mem. S. A. It., 72, 299
- Pols O. R., Tout C. A., Eggleton P. P., Han Z., 1995, MNRAS, 274, 964
- Stancliffe R. J., 2005, PhD thesis, University of Cambridge
- Stancliffe R. J., Izzard R. G., Tout C. A., Pols O. R., 2004, Memorie della Societa Astronomica Italiana, 75, 670
- Stancliffe R. J., Lugaro M. A., Ugalde C., Tout C. A., Görres J., Wiescher M., 2005, MNRAS, 360, 375
- Stancliffe R. J., Tout C. A., Pols O. R., 2004, MNRAS, 352, 984
- Straniero O., Limongi M., Chieffi A., Dominguez I., Busso M., Gallino R., 2000, Mem. S. A. It., 71, 719